

1 PROCESSOR STATE AWARE INTERRUPTS FROM PERIPHERALS

2 FIELD OF THE INVENTION

3 This invention is related to interrupt processing in
4 computer systems. More particularly, it relates to
5 interrupt processing in a manner so as to reduce power
6 consumption and conserve processing resources.

7 BACKGROUND OF THE INVENTION

8 Currently the peripheral units in a system, activate
9 their interrupt lines whenever they need attention from
10 the processor without any concern for what the
11 processor may be doing. In a low power system, the
12 processor may be in either an active state or a low
13 power sleep state. It usually takes a finite number of
14 cycles for a processor to transition into the low power
15 sleep state from the active state, and similarly a
16 finite number of cycles to transition from the low
17 power sleep state to the active state.

18 When a peripheral activates its interrupt line, the
19 processor transitions from the low power sleep state
20 into an active state, if it was sleeping, to respond
21 to the interrupt. As an example, there can be two
22 peripherals A and B, each with a separate interrupt
23 line. When the processor is in the sleep state, if
24 peripheral A needs attention, the voltage on its
25 interrupt line is changed. The processor will then
26 come out of the sleep state into the active state to

1 handle interrupt A and after it is done, the processor
2 goes back to sleep. A short while later peripheral B
3 activates its interrupt line and the processor repeats
4 the wake-up sequence to handle interrupt B. This is
5 wasteful in terms of power consumed and system
6 resources used.

7 In general many peripherals have some amount of
8 buffering which they can use to implement a certain
9 level of slack with respect to when they need to
10 interrupt the processor. For example, a serial
11 interface may have a sixteen byte first in first out
12 (fifo) memory to hold incoming characters. The serial
13 interface may be configured to interrupt the processor
14 as soon as one character has come in, or may be
15 configured to interrupt the processor when the fifo
16 memory is half full. Other options may be possible as
17 well. For example an interrupt can be activated after
18 one character is received if no character is
19 subsequently received for two character times.

20 Similarly, disk requests from a processor typically go
21 into a queue. The processor sets up several disk writes
22 and reads and triggers the disk controller. The disk
23 controller processes elements from the queue and can
24 interrupt the processor at different times, after each
25 successful operation, when the queue is half empty or
26 when the queue is fully empty.

1 Yet another example relates to networking. Similar to
2 disk operations, network transfers can also be queued.
3 The network interface has the option of interrupting
4 the processor at different thresholds.

5 In some cases the changing of thresholds may affect the
6 correctness or smooth operation of the system. For
7 instance if the serial interface delays the delivery of
8 incoming bytes to the processor, the processor may not
9 acknowledge receipt of the bytes and thereby prevent
10 the transfer of subsequent bytes on the same serial
11 line. However in many other cases, it is acceptable to
12 modify the thresholds where the peripherals need to
13 signal the processor. The setting of these thresholds
14 is often driven by optimizing some metric such as user
15 response time or total throughput depending on whether
16 the machine is to be used as an interactive workstation
17 or a server.

18 SUMMARY OF THE INVENTION

19 It is an object of the present invention to service
20 interrupts from peripherals in a manner that is
21 conservative of energy and system resources.

22 It is a further object of the invention to synchronize
23 the servicing of interrupts from peripheral devices.

1 It is yet another object of the invention to
2 efficiently service interrupts from peripherals by a
3 processor having a number of distinct processor
4 activity states.

5

6 The present invention is based in part on the
7 recognition that if the interrupts could be
8 synchronized in some way, so that requests from
9 peripherals A and B both can be handled in a single
10 wake-up transition, the total energy consumed is lower.

11 In accordance with the invention, thresholds are
12 automatically adjusted based on the current state of
13 the processor. In particular, in a preferred
14 embodiment, the processor provides an output signal,
15 possibly on one or more lines, that is indicative of
16 the state the processor is in (for example, an active
17 state or a sleep state). The peripheral units are
18 connected to this (or these) signal line. The
19 peripherals monitor this signal and their interrupt
20 thresholds are varied to be low when the processor is
21 active and to be high when the processor is asleep. In
22 essence what this does is cause the peripherals to
23 delay their respective interrupts when the processor is
24 asleep.

25 When the processor is asleep, all peripherals hold off
26 their interrupts until one of them hits a high urgency
27 threshold. This peripheral interrupts the processor
28 waking up the processor. Once the processor is awake

1 all other peripherals activate their interrupts if
2 their low threshold has been crossed, effectively
3 causing the processor to handle all of the peripherals
4 in one wake up sequence.

5 This mechanism can be easily generalized to the case
6 where the processor supports multiple low power levels,
7 such as idle, sleep, or deep sleep. When there are more
8 states, the processor needs to put out multiple bits of
9 output so that the processor state can be encoded. For
10 instance if there are four states, in a hardware
11 embodiment of the invention, two wires are needed.

12 In general, peripherals are able to determine how
13 urgent the need is for processor attention.
14 Peripherals also monitor the processor to see what
15 state it is in. The deeper the sleep state of the
16 processor, the longer the peripherals hold off their
17 interrupt i.e., they wait until their level of urgency
18 is very high.

19 If a peripheral in a state of high urgency interrupts
20 the processor and wakes it up, all peripherals which
21 are at lower levels of urgency raise their interrupt
22 levels asking for processor attention. This mechanism
23 automatically aligns all interrupts, thus enabling the
24 processor to do a great deal all in one sweep, rather
25 than waking up repeatedly and going into deep sleep. In
26 other words, this mechanism is automatically self
27 synchronizing in that an awake processor automatically

1 triggers all peripherals that may need service in the
2 near future, to request service and thereby clear their
3 work queues. In addition, once all the peripherals have
4 been serviced and the processor goes to sleep, the
5 peripherals automatically hold off on their interrupts
6 until one of them reaches a high work threshold (high
7 state of urgency).

8 **BRIEF DESCRIPTION OF THE DRAWINGS**

9 These and other aspects, features, and advantages of
10 the present invention will become apparent upon further
11 consideration of the following detailed description of
12 the invention when read in conjunction with the drawing
13 figures, in which:

14 Fig. 1 is a block diagram of a prior art system.

15 Fig. 2 is a block diagram of a system in accordance
16 with the invention.

17 Fig. 3 is an exemplary flow chart of the operation of
18 the system of Fig. 2.

19 Fig. 4 is an exemplary timing diagram of the operation
20 of the system of Fig. 2.

1 DESCRIPTION OF THE INVENTION

2 Variations described for the present invention can
3 be realized in any combination desirable for each
4 particular application. Thus particular limitations,
5 and/or embodiment enhancements described herein, which
6 may have particular advantages to the particular
7 application need not be used for all applications.
8 Also, it should be realized that not all limitations
9 need be implemented in methods, systems and/or
10 apparatus including one or more concepts of the present
11 invention.

12 Referring to Fig. 1, prior art computer system 10 has a
13 main processor 12 that has multiple interrupt lines 14.
14 Each interrupt line is assigned to a particular
15 peripheral interface 16. A shared interrupt line 18 is
16 shared amongst multiple peripheral interfaces 20. Each
17 peripheral interface has connections to the external
18 world I/O devices such as keyboard, mouse, network,
19 disk, etc.

20 Referring to Fig. 2, in accordance with the invention,
21 the structure of Fig. 1 is enhanced by adding one (or
22 more) lines, as represented by 22, that are output from
23 the processor and that indicate its current state. If
24 the processor can be in more than two states, one line
25 or wire may be inadequate. If the processor can be in
26 the states of "Active", "Idle", "Sleep" two lines
27 having only binary outputs (a "1" or a "0") thereon are

1 needed to indicate one of three possible states. These
2 lines are connected to all the peripheral interfaces 16
3 and 20, thus supplying information to the interfaces to
4 determine the current state of the processor at any
5 point in time by determining the potentials on these
6 lines.

7 Referring to Fig. 3, each peripheral interface goes
8 through the flow chart that is presented. Normally the
9 peripheral interface is waiting 300 for something to
10 happen. If it sees an external I/O event 302, it first
11 enqueues the event 304 and checks the current processor
12 state 306. As explained above, each I/O event has some
13 effect on or internally changes the state of the
14 interface to some level of criticality C0, C1, C2,
15 etc., where C0 is less than C1, which is in turn less
16 than C2. Based on the current processor state
17 determined at 306, the peripheral unit compares its
18 internal level of criticality against different
19 thresholds C0, C1 or C2 as appropriate. If the
20 processor is active, then any criticality greater than
21 C0, at 308, will activate an interrupt for that
22 processor at 310. If the processor is in an idle state,
23 then any criticality greater than C1, at 312, will
24 activate an interrupt for the processor at 310. If the
25 processor is in a sleep state, then any criticality
26 greater than C2, at 314, will activate an interrupt for
27 the processor at 310. In short, if the level of
28 criticality is higher than the appropriate threshold,
29 the peripheral interface activates its interrupt line

1 asking for the processor to service the interface. If
2 the criticality is lower than the threshold, the
3 interface does nothing and waiting 300 continues. If
4 the processor changes its state 316 (perhaps due to
5 some other peripheral interface interrupting the
6 processor), the peripheral interface in question
7 detects this and then again runs the threshold checker
8 at 306. The threshold of interest may have become lower
9 due to the processor being in a more "awake" state. If
10 this is the case, the peripheral unit activates its
11 interrupt line.

12 Fig. 4 shows a sample runtime behavior. Going forward
13 in time, from left to right, the processor transitions
14 from active to idle and finally to sleep since it has
15 nothing to do. When the processor is in the sleep
16 state, external events occur on Peripheral Interface 1
17 that raise its level of criticality gradually, but the
18 level of criticality does not exceed the Sleep state
19 threshold (C2) for Peripheral Interface 1. An event
20 occurs on P5 that raises its criticality level but this
21 is still lower than C2 for P5. Finally another event
22 occurs on Peripheral Interface 5 that puts it above its
23 threshold C2 causing it to activate its interrupt line.
24 The processor immediately wakes up and services
25 Peripheral Interface 5. As it wakes up, the processor's
26 state goes to "Active", causing P1 to reevaluate. Now
27 since its level of criticality is higher than C0, it
28 activates its interrupt line. After the processor has
29 completed servicing P5 it services Peripheral Interface

1 1. As each peripheral unit is serviced, its level of
2 criticality drops to zero. Finally the processor has
3 completed all its activity and it drops to the idle
4 state for some, generally predetermined, period of
5 time. After a time-out period has elapsed the processor
6 drops down to the even lower powered sleep state.

7
8 The net effect of all of these changes is that the
9 processor is awakened less frequently from its lowest
10 power state and can save more energy because of that.
11 It also ensures that when the processor wakes up, it
12 deals with all the peripherals in quick succession
13 thereby amortizing the cost of state transitions.

14 While an implementation of the invention has been shown
15 which uses one or more signal lines, it will understood
16 by one skilled in the art that the activity state of
17 the processor may also be supplied to the peripherals
18 by sending specifically coded digital information along
19 one or more existing communication lines between the
20 processor and the peripheral. For example at least one
21 output word may be generated by the processor and
22 communicated to the peripherals, which is indicative of
23 the activity state of the processor. Thus, the
24 invention may be implemented without adding additional
25 hardware signal outputs from the processor. One
26 possible approach is to modify existing peripheral
27 firmware to be responsive to digital words from the
28 processor indicative of the activity state of the
29 processor, and to internally store the processor state
30 and any changes to the processor state, in responsive

1 to the digital word on the existing communication
2 lines.

3 The present invention can be realised in hardware,
4 software, or a combination of hardware and software.
5 Any kind of computer system - or other apparatus
6 adapted for carrying out the methods and/or functions
7 described herein - is suitable. A typical combination
8 of hardware and software could be a general purpose
9 computer system with a computer program that, when
10 being loaded and executed, controls the computer system
11 such that it carries out the methods described herein.
12 The present invention can also be embedded in a
13 computer program product, which comprises all the
14 features enabling the implementation of the methods
15 described herein, and which - when loaded in a computer
16 system - is able to carry out these methods.

17 Computer program means or computer program in the
18 present context include any expression, in any
19 language, code or notation, of a set of instructions
20 intended to cause a system having an information
21 processing capability to perform a particular function
22 either directly or after conversion to another
23 language, code or notation, and/or reproduction in a
24 different material form.

25 Thus the invention includes an article of manufacture
26 which comprises a computer usable medium having
27 computer readable program code means embodied therein

1 for causing a function described above. The computer
2 readable program code means in the article of
3 manufacture comprises computer readable program code
4 means for causing a computer to effect the steps of a
5 method of this invention. Similarly, the present
6 invention may be implemented as a computer program
7 product comprising a computer usable medium having
8 computer readable program code means embodied therein
9 for causing a function described above. The computer
10 readable program code means in the computer program
11 product comprising computer readable program code means
12 for causing a computer to effect one or more functions
13 of this invention. Furthermore, the present invention
14 may be implemented as a program storage device readable
15 by machine, tangibly embodying a program of
16 instructions executable by the machine to perform
17 method steps for causing one or more functions of this
18 invention.

19 It is noted that the foregoing has outlined some of the
20 more pertinent objects and embodiments of the present
21 invention. The concepts of this invention may be used
22 for many applications. Thus, although the description
23 is made for particular arrangements and methods, the
24 intent and concept of the invention is suitable and
25 applicable to other arrangements and applications. It
26 will be clear to those skilled in the art that other
27 modifications to the disclosed embodiments can be
28 effected without departing from the spirit and scope of
29 the invention. The described embodiments ought to be
30 construed to be merely illustrative of some of the more

1 prominent features and applications of the invention.
2 Other beneficial results can be realized by applying
3 the disclosed invention in a different manner or
4 modifying the invention in ways known to those familiar
5 with the art. Thus, it should be understood that the
6 embodiments has been provided as an example and not as
7 a limitation. The scope of the invention is defined by
8 the appended claims.